



(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention
of the grant of the patent:
17.11.1999 Bulletin 1999/46

(21) Application number: **91900886.2**

(22) Date of filing: **21.12.1990**

(51) Int Cl.⁶: **G06F 15/80**

(86) International application number:
PCT/GB90/02006

(87) International publication number:
WO 91/10196 (11.07.1991 Gazette 1991/15)

(54) **NEURAL NETWORKS**
NEURONALE NETZWERKE
RESEAUX NEURONAUX

(84) Designated Contracting States:
DE DK ES FR GB IT NL SE

(30) Priority: **22.12.1989 GB 8929146**

(43) Date of publication of application:
07.10.1992 Bulletin 1992/41

(73) Proprietor: **BRITISH TELECOMMUNICATIONS**
public limited company
London EC1A 7AJ (GB)

(72) Inventors:
• **NIGHTINGALE, Charles**
Felixstowe, Suffolk IP11 7JL (GB)
• **WYARD, Peter Joseph**
Ipswich, Suffolk IP5 7UW (GB)

(74) Representative:
Roberts, Simon Christopher et al
BT Group Legal Services,
Intellectual Property Department,
8th Floor, Holborn Centre
120 Holborn
London, EC1N 2TE (GB)

(56) References cited:
EP-A- 0 327 817

- **IEEE International Conference on Neural Networks, San Diego, California, 24-27 July 1988, S. Nolfi et al.: "Learning to understand sentences in a connectionist network", pages II-215-II-219**
- **IEEE International Conference on Neural Networks, San Diego, California, 24-27 July 1988, M.F. Tenorio et al.: "Adaptive networks as a model for human speech development", pages 235-242**
- **Neural Networks from Models to Applications, Paris, 1988, Ekeberg et al.: "Automatic generation of internal representations in a probabilistic artificial neural network", pages 178-186**
- **IJCNN International Joint Conference on Neural Networks, Washington, 19-22 June 1989, T. Matsuoka et al.: "Syllable recognition using integrated neural networks", pages 251-258**
- **Neuro-Nimes'89, Nimes, 13-16 November 1989, P.-P. Meiler: "Garbled text string recognition with a spatio-temporal pattern recognition (SPR) neural network", pages 279-292**

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

[0001] This invention relates to (artificial) neural networks (in other words, to parallel processing apparatus comprising or emulating a plurality of simple, interconnected, neural processors, or to apparatus arranged to emulate parallel processing of this kind) and particularly, but not exclusively, to their use in pattern recognition problems such as speech recognition, text-to-speech conversion, natural language translation and video scene recognition.

[0002] Referring to Figure 1, one type of generalised neural net known in the art comprises a plurality of input nodes 1a, 1b, 1c to which an input data sequence is applied from an input means (not shown), and a plurality of output nodes 2a, 2b, 2c, each of which produces a respective net output signal indicating that the input data sequence satisfied a predetermined criterion (for example, a particular word or sentence is recognised or an image corresponding to a particular object is recognised). Each output node is connected to one or more nodes in the layer below (the input layer) by a corresponding connection including a weight 3a-3i which scales the output of those nodes by a weight factor to provide an input to the node in the layer above (the output layer). Each node output generally also includes a non-linear (compression) stage (not shown).

[0003] In many such nets, further intermediate inner or 'hidden' layers are included, which receive inputs from a layer below and generate outputs for a layer above. The output of a node in general is a function of its weighted inputs; typically the function is the sum of these inputs, with the subsequent non-linear compression mentioned above. One example of such a net is the well known Multi-Layer-Perceptron (MLP).

[0004] Such nets are trained in a training phase by inputting training data sequences which are known to satisfy predetermined criteria, and iteratively modifying the weight values connecting the layers until the net outputs approximate the desired indications of such criteria.

[0005] Having been trained on a range of training data, it is then found that such trained networks can operate upon real-world data to perform various processing and recognition tasks.

[0006] Since the revival of interest in neural nets in recent years much attention has focussed on nets in which processing is unequivocally parallel and distributed, (Rumelhart 1986 [8]) and which have recently proved to be admirably suited to tackling problems in signal processing eg (Lynch & Rayner 1989) pattern recognition eg (Hutchinson & Welsh 1989) (Woodland & Smythe 1990) and robotic control eg (Saerens & Soquet 1989). Some attention has also been paid to problems which cannot be seen as signal processing, and in particular various methods of applying neural nets to natural language have been described, from (Rumelhart 1986 [9]) and (McClelland & Kawamoto 1986) through to recent papers and reports (Sharkey 1989), (Weber 1989) and (Jagota & Jajubowitz 1989). A difficulty in these cases is how to present inputs to the net. If unlimited data such as text is to be processed by a neural net of these kinds, either it must be input as some set of lower level features - letters or microfeatures as described in eg (Rumelhart *et al* 1986 [10]), - or if whole words or larger features are to be used the number of input nodes must be very great. In the latter case, too, some retreat from the pure concept of parallel distributed processing must be accepted, since each word can be seen as locally stored.

[0007] In other words, the choice is typically between using too few nodes (in which case the network may not train well if features chosen are inappropriate) or too many (in which case the network is tending to act as a simple look up store).

[0008] Another problem is that a very large number of iterations can be required for convergence in training, which can consequently be slow and laborious.

[0009] In their paper entitled "Learning to understand sentences in a connectionist network", published in the proceedings of the IEEE International Conference on Neural Networks, San Diego, 24-27 July 1988, pages II 215 to II 219, Nolfi and Paris; describe a "Jordan Architecture" net which is trained by back propagation. The net is a kind of multi-layer perceptron in which there are input units, output units and hidden units. Associated with each hidden unit is a corresponding memory unit. Each memory unit makes a temporary copy of each state of its associated hidden unit and then supplies this copy to the hidden unit in the next cycle (when the system processes the next stimulus).

[0010] The memory units only store information temporarily. The information stored in the memory units does not appear to correspond to the "new features" specified herein. The information stored in the memory units is not used to modify the input layer in any way.

[0011] In another paper at the same conference, at pages II 235-242, Tenorio *et al* discuss the NETTalk system applied to Spanish and English. In paragraph 5.2 of that paper they discuss the effects of using networks having different numbers of hidden units. Unsurprisingly, when the network has very many hidden units (there being at least as many hidden units as there are training patterns), rather than only a few, there is a dramatic change in the performance in the back propagation algorithm. With many hidden units the network can of course operate as what is effectively a look up table. There is no suggestion either that there is an optimum number of hidden units or that the number of hidden units in a particular network should be altered dynamically or in any other way.

[0012] Ekeberg, in a paper entitled "Automatic generation of internal representations in a probabilistic artificial neural network", published in "Neural Networks from Models to Applications", I.D.S.E.T. Paris 1989, at pages 178 to 186,

considers adding layers and features to what is initially a single layer feedback perceptron type network. Higher level features code for suitable combinations of simultaneous input unit activity. The higher level features are present in a separate layer, communicating through connections with the input/output layer. Initially the internal layer contains one unit for each in/out unit, that is the internal code is initially the same as the external. During training, the internal layer is gradually transformed by replacing existing units with units coding higher order co-activity in the In/Out Layer. Ekeberg describes how the appropriate internal units are chosen: he selects the two internal units with the highest interdependency and replaces them with three more specific ones, one being active when both of the old ones were, and each of the other being active when only one was active. After such a replacement, the sample patterns are scanned again to get estimates of the new probabilities involved. This process of replacing two internal units by three is repeated until the task is solved. In the worst case, when no useful regularities are detected, so-called "Grandmother cells" corresponding to the individual training patterns will develop. Thus the network grows into what is effectively a look-up table, which of course is very memory intensive. Ekeberg states that in the normal case, however, a successful representation emerges much sooner.

[0013] Ekeberg does not suggest either the possibility or desirability of limiting the replacement of units for the normal case. Ekeberg is also silent as to whether or how the creation of "grandmother cells" can be limited or inhibited with or without any deterioration of the "normal case" performance.

[0014] In EP-A-0327817 there is described an associative pattern conversion system in which, during training, connection weights are adjusted between pre-determined, fixed maximum and minimum values. The maxima and minima are fixed in advance, and preferably with only a small range between them, in order that a simple circuit can be used. When a weight reaches its pre-determined maximum or minimum, it is said to be saturated. The weight modification function is a monotonically decreasing function. There is no suggestion of adding internal nodes or modifying the input layer to be responsive to new features derived from higher level features.

[0015] According to a first aspect the invention provides a trainable artificial neural network as claimed in Claim 1.

[0016] The net uses extra memory (or nodes) to deal with difficult training tasks, in preference to long training programs. Even so the memory used is not excessive, as restraints are preferably placed upon the net's propensity to create new nodes.

[0017] According to a further aspect the invention provides a method of training an artificial neural network as claimed in Claim 10.

[0018] Other aspects and embodiments of the invention are as recited in the appended claims, or as described hereafter.

[0019] The invention will now be described by way of example only, with reference to the accompanying drawings in which:

Figure 1 shows schematically a general (prior art) neural net,

Figures 2a - 2d show schematically the structure of a net according to one aspect of the invention during subsequent stages of training,

Figure 3 shows schematically a general (prior art) output or intermediate node of a neural net,

Figure 4 shows schematically a method of training the net of Figures 2a - 2d,

Figure 5 shows schematically a weight modification function according to one aspect of the invention,

Figure 6 shows schematically an embodiment of the invention for grammatically checking,

Figures 7a and 7b illustrate a net according to a further aspect of the invention during subsequent stages of training,

Figure 8 illustrates a net according to this embodiment trained to solve one problem, and

Figure 9 illustrates a net according to this embodiment trained to perform the XOR logic function.

[0020] Referring to the drawings, in a simple example, when the net starts its training phase it may consist of only a single output node which fires in response to any input. As it receives new and unfamiliar inputs, the net instantiates new nodes as necessary each responsive to a feature found in the input data. Therefore after the simple form of the invention has been running for some time it will consist of a layer of input nodes, connected by a set of connections to a single layer of output nodes, as shown in Figure 2a. The connections from input to output do not exist between every input/output pair, but each connection possesses a weight. Input signals may comprise, for example, text, speech or video data. In an example where the input consists of text entities, if words are detected as the lowest level features, then a complete training phase input data sequence could be a phrase or sentence. Each node in the input corresponds (i.e. produces an output in response) to some feature of an input sequence that has been applied to the net previously, and when a new input data sequence is applied which includes some of these low-level features, these nodes will become excited. Any data representing unknown low-level features contained in the input thus play no part in the determination of the output of the net. Such data are however retained in short term memory during the subsequent cycles of connection-weight modification, and are ultimately connected to some appropriate output so as to form a new input node. Each output node is a typical neural processor, as shown in Figure 3 in which the excitation is given by Eqn 1

$$y = \sum_{i=1}^n w_i x_i \quad (1)$$

[0021] No nonlinear activation function as such appears here, but a competitive algorithm, which ensures that a notional thresholding operation occurs which is automatically set to cut off all but the strongest firing output cell, fulfills the same function (a nonlinear function could, however, be used).

[0022] The algorithm works, in a simple single-layer mode, by following the flow-diagram shown in Figure 4 which will be described below.

[0023] Referring to Figure 4 the algorithm works as follows:-

[0024] Ready1: At position 6 in the interaction cycle the system is in its start state and the user has the choice of proceeding via 1, 2 or 3.

[0025] Interrogate: At position 1 of the interaction cycle assume the net configuration is as shown in Figure 2a.

[0026] A new input is given to the net which excites some of the cells in the input layer, say a2, a4 and a6, and instantiates (i.e. creates by storing feature data for) two new cells a8 and a9. The cells (a2, a4, a6, a8, a9) which together store the entire input data sequence are known in the net as the current short term memory (CSTM) and are conveniently provided in RAM. Several cells in the output layer (b1, b3 and b4) are excited by the firing of the input cells via appropriate existing interconnections eg (a2-b3, a4-b4, a6-b1 and a6-b3). The output layer behaves exactly like a cluster in competitive learning (Rumelhart & Zipser 1986) in that one cell then dominates all the others and inhibits all outputs except its own. (In principle this could be done by fully connecting the output layer with inhibitory connections, and using some feedback system to stabilise the net when only a single cell remains firing. In practice this is cumbersome, and does not contribute to either the ease of implementation or the understanding of the net behaviour, so a simple global function which identifies the cell which fires most strongly is applied, for the purpose of simplification). Suppose it is b1 that dominates in our example. This cell then fires and produces a certain action in the net. This action will normally mean printing an output, but could equally lead to other things. Unlike known competitive learning algorithms there is no weight modification at this stage. The program now returns to Position 6. The supervisory aspect of the net is brought into play at this stage. The supervisory program, or user, decides whether the net has made a correct or incorrect response. In the event that the response is correct the net can either be left unmodified, or it can be "rewarded" (i.e. modified to encourage this response). If the response is considered incorrect then the net should be punished, where punishment is a weight modification process which makes the incorrect response less likely for the given input. In the event that the response is considered neither correct nor incorrect then the net can be left unmodified and further inputs can be tried.

[0027] Reward: At position 2 assume the net is as shown in the example in Figure 2b and has output the correct response. The user rewards the net, and the connections are modified in the following way: Connections to the excited output node which are not carrying a signal are weakened eg a1 - b1. Connections to the newly instantiated nodes are made eg a8 - b1, a9 - b1.

[0028] Connections to the excited output node which are carrying a signal are strengthened eg a6 - b1.

[0029] Other connections are unaltered.

[0030] The algorithm then returns to 6.

[0031] Punish: At position 3 in the interaction cycle, the net as shown in Figure 2c has produced an incorrect response, so the user punishes the net, and the connections are modified in the following way:

[0032] Connections to excited output node which are not carrying a signal are strengthened eg a1 - b1.

[0033] Connections to excited output node which are carrying a signal are weakened eg a6 - b1.

[0034] Other connections are left unaltered.

[0035] The cycle then proceeds to Position 5.

[0036] Ready2: At this stage it is possible either to proceed to position 4 or return to Position 6.

[0037] Teach: At Position 4 the user gives the expected output for the net to learn: for example assume that b3 is the expected output as shown in Figure 2d.

[0038] Teach proceeds as for reward above - making connections between all CSTM input cells and expected output cell which do not already exist eg a4 - b3, a8 - b3, a9 - b3, a6 - b3.

[0039] Other connections to given cell are weakened, e.g. a7 - b3.

[0040] The program then returns to position 6.

Weight Modification

[0041] The method of altering the weights is a modification of the method used in competitive learning, (Rumelhart & Zipser 1989) except that instead of strengthening connections in such a way as to reinforce the existing tendencies of the net as is done in competitive learning, the procedure is controlled so that only desired responses are strengthened and undesired responses are weakened.

[0042] In competitive learning it is usual to normalise the weights, but in the DTN a different strategy is adopted, as shown in Equation 2

$$W_{new} = \left[1 + \frac{\mu \delta}{1 + (\mu \delta W_{old})^4} \right] W_{old} \quad (2)$$

where $\delta = +1$ causes strengthening, and $\delta = -1$ causes weakening μ governs the position of the maximum.

[0043] Figure 5 shows graphs of this weight modification function, for strengthening and weakening. The effect of this weight modification function is that when first instantiated weights are easy to alter (because the weights are around the minimum of the function), but when a lot of weakening or strengthening has taken place the strengths tend to saturate, and do not easily increase or decrease. This prevents weights from becoming too dominant, but ensures that if a weight has been constantly strengthened it is not easily weakened. Unlearning is therefore difficult, though possible. This weight modification method is also applicable to known types of neural network, for example MLP networks.

[0044] As it stands the net described above is capable of useful associations, and was initially used as a top-down contexter in a scene understanding program. In, for example, a scene in which ships and water has been detected, an overall candidate for the type of scene (a harbour, say) could be determined using the net, and any large objects including vertical and oblique lines could become candidate cranes to be tested using a bottom up process. The net has no sense of structure of the input, at this stage, so that it cannot distinguish between, for example, (boat water crane) and (crane boat water). This is not a fatal weakness as far as the contexting purpose is concerned, but for other experimental applications, such as limited domain translation, the inability to distinguish between, (what time is the next train from London to Harwich) and (what time is the next train from Harwich to London) would have been unacceptable.

[0045] In embodiments for language processing and similar order-dependent input data, the net therefore needs some means of retaining the order dependence.

[0046] In one embodiment this flexibility is achieved by creating new input nodes each responsive to more than one feature. Thus, when a sequence of training data is input, it is temporarily stored in an input buffer. As in the simple embodiment, known features (for example, words of a sentence) cause certain input nodes to fire and certain new nodes, responsive to new features, may be created, as discussed above. Then, in this embodiment of the invention, order information is captured by forming further new nodes each responsive to a plurality of features in the input data sequence; preferably each plurality is a contiguous subsequence.

[0047] For example, if the training data sequence is a sentence, this embodiment captures the order information by forming contiguous pairs, triples, 4-tuples etc. of preterminals (or words) and storing each of these in an input node.

[0048] In the sentence "lazy cats sleep in the hot sun", for example, the pairs "lazy", "cats", "cats sleep", "sleep in" etc; the triples "lazy cats sleep", "cats sleep in", "sleep in the" etc, the 4-tuples "lazy cats sleep in", "cats sleep in the" etc could be formed, and stored in short term memory, and corresponding nodes may be formed for some for some or all. Such nodes will preferably be removed from the net if not encountered subsequently during training.

[0049] Figure 6 shows a portion of a net according to this embodiment after training on the problem of deciding whether a given sentence is grammatical. (In a typical case there could be 30 or 40 nodes in the input layer after training on the grammaticality problem, or thousands of nodes after training on the problem of language translation).

[0050] The algorithm or mode of operation of this embodiment during training will now be outlined, taking the grammaticality problem as an example. The training file consists of a set of positive and negative sentences of preterminals, each one followed by its correct classification *yes* or *no* (the "desired response" in the program). One presentation of a sentence leads to the following sequence of operations (the "main cycle"):

1) The training data (sentence) string is read from the file

2) All possible tuples are generated in short term storage (between the minimum and maximum tuple lengths specified by the supervisory program or user)

3) Some of these tuples are selected for incorporation into the input layer. The selection is discussed below in detail, but is essentially pseudo-random: the probability of selection depends on several factors to be discussed below, and also upon the tuple length (the longer the tuple, the less likely to be selected)

4) Each input node matching any of the generated tuples is activated (i.e. fired), and the activations of the output nodes are calculated by a simple weighted sum of the active input nodes

5) All tuples selected in step 3 are added to the input layer to form new nodes (if they are not already present)

6) The input layer is pruned of nodes which have not been very "useful", i.e. infrequently activated over the preceding cycles

7) The most active output node is found, and designated "current output_winner"

8) The desired response is read from the file. If this is the same as the current output winner, the latter is simply marked "desired_response_node" as well. If it is not the same, then this node is added to the output layer if it does not already exist, and it is marked as "desired_response_node"

9) Links are created if necessary between all active input nodes (including the ones just added) and the desired response node

10) Weights of links are adjusted according to the learning method described above (with $\mu = 1$, for example).

[0051] As more and more sentences are presented during training, the input layer first grows and then reaches a state of dynamic equilibrium where nodes are being lost as fast as they are being added. The training regime can be either "brute force" where one simply runs through the entire training set time after time until performance is judged satisfactory, or "incremental" where sentences from the complete training set are only added to the current training set when performance on the latter is satisfactory. (What is "satisfactory" is something to be discovered experimentally, with reference to performance on the test set. It was usually found best to demand only 90% performance on the training set rather than 100%, since this gave considerably quicker training and hardly degraded performance on the test set.)

[0052] As stated above, the training can be either "brute force" or "incremental". Brute force training is the standard mode of training neural nets (particularly MLPs) where one simply cycles repeatedly through the entire training set. Although this appears to work perfectly well for MLPs, it does not work at all well for this embodiment with training sets larger than toy size (the difference is thought to be linked to the different weight adjustment methods used). Even a training set of 40 strings could not always be learnt easily using brute force training.

[0053] It is therefore preferred to use an "incremental" training method, in which the network is trained on a first set of training data sequences, and after training is complete only then the set is expanded. Training data sequences may be added to the set one at a time or in steps of several sequences. The net retrains after each expansion of the training data sequence set. In general, the larger the incremental stepsize the less the total number of presentations required for learning a training set, but if the stepsize was made too big then learning would become slower again. The optimum stepsize must be found by trial and error, since it is data dependent.

[0054] This incremental training method is also applicable to other embodiments of the invention, and to other types of neural network (eg MLP networks) but is preferably used in combination with the above weight adjustment method.

[0055] As discussed above in step 3, not all tuples of an input data sequence will always be stored because in many domains the number of input nodes would grow too big for practical computation if this were done. Some means of selecting for storage only some of the tuples is required. The simplest solution would have been to store every second or every third tuple encountered during training, say, but this is usually too crude a criterion. A probabilistic method of storage was therefore developed, with the probability varying during training in the way described below.

[0056] The probability of incorporating a tuple of length n into the input layer depends on p (a global variable to be described below) and n . At present, this embodiment uses a simple power law: the probability of incorporating a tuple of length n is p^n . Thus if $p = 0.8$, the probability of creating a singleton is 0.8, of creating a pair is $0.8^2 = 0.64$, of creating a triple is $0.8^3 = 0.512$, etc., up to the maximum tuple length specified. More sophisticated probability distributions could be considered. The rationale behind the power law is that the longer tuples are in a sense more "specialised" and not so generally useful as the shorter ones for selecting the desired output node.

[0057] It turned out that for simple problems on grammaticality and translation that it was simplest to keep p fixed at 1.0, i.e. to store *all* tuples encountered during training, up to the maximum tuple length. This was computationally feasible as long as one only stored short tuples - up to a maximum length of, say, 3. The advantage of this approach

is that one does not have to worry about very useful tuples being thrown away due to chance. However, probabilistic storage will be required in order to learn larger data sets containing more tuples.

[0058] P will also depend on global net parameters. A first such parameter may be termed "complacency". Complacency is related to the length of the current sequence of correct answers. If the net is responding correctly we do not wish to add many new nodes. This general principle is also applicable to all other embodiments of the invention. In this embodiment complacency causes p to decay exponentially as the correct sequence grows in length, then return to its original value when a wrong answer is produced.

[0059] A further parameter, which may be termed "experience", is related to the total number of nodes in the input layer. We do not wish to add nodes without limit when the net is already big, so experience causes p to decay exponentially with respect to the number of input nodes. Complacency is useful in a situation where the net has been exposed to part of a very large data set during incremental training, which was already sufficient to make the net perform at a very high level in the particular problem area. The idea is that the rest of the data set will not "clog up" the input layer by causing the addition of many more superfluous input nodes. However, in many types of situation, the training data set is not large enough to make the use of "complacency" essential. The use of "experience" could potentially cause problems if the net had already grown large, thus making tuple storage less likely, but had not yet learnt a training set. In such a case it might never learn the training set. It is therefore important to ensure that p only decays extremely slowly with experience.

[0060] This embodiment of the invention is, as stated above, useful in language processing problems (although it may also be used for other types of data having a feature order dependency) and may, for example, be used to train a phrasebook-type language translation device.

[0061] An advantage of such a neural net approach to translation would be that a net could be trained by anyone who spoke both source and target language, with an appropriate method of data collection, and would not require skilled programming by language experts. In cases where it is economic to spend a lot of resources on a powerful translation program, for example in a Japanese-English context, it would be feasible to implement an elaborate classical language translation algorithm. In cases where languages spoken by smaller groups of less economic strength are involved, it may well be advantageous to use a neural net aided translator. Any such net could, of course, produce the intended output phrase in the source language as well as the target language so that serious mistakes could be filtered out.

[0062] Another simple method for giving the inputs some ordering information, in which it was possible to associate an absolute position with each symbol (for certain symbol sets) is as follows.

[0063] For example if the multiple inputs L O O L L O L L were input, the net included a preprocessor which converted them into L1 02 03 L4 L5 06 L7 L8, which can then be treated as a single binary input (10011011).

[0064] This has the drawback of giving the net a notional rigidity - since it makes sense to then consider inputs of a fixed length - but it is no more rigid than a multi-layer-perceptron for example, although a more flexible way of introducing order is desirable.

[0065] The use of the above types of input layer are found effective for certain types of problem. For higher level problems, use of hidden or intermediate layers of nodes between the input and output layers is often an effective solution. Although this could be of a conventional kind, networks according to a further embodiment of the invention have the additional capability of creating (additional) internal nodes (i.e. in an intermediate layer) to cope with what appear to a single layer net to be contradictory inputs. If such nodes were created too frequently, and for cases where a single-layer net could cope if enough training were given, then the net would soon be swamped with a large number of unnecessary internal nodes, and what would have been created would be equivalent to nothing more than a very large memory. If on the other hand such cell creation is usually only performed when necessary, to cope with complex data, then a net with such a capability can learn difficult problems without using excessively large amounts of memory. Mechanisms whereby such extra nodes can be introduced sparingly, but appropriately will now be discussed.

[0066] A global parameter equivalent to the "complacency" parameter above is employed (in the following, this version of the complacency parameter will be termed 'Cell Creation Excitation' or 'CCE').

[0067] CCE is conveniently related to the number of punishments the network has historically received during training compared to the number of interrogations. If the network has given many correct answers and hence received few punishments the network will have a low CCE and will be unlikely to create, or instantiate, any new cells. In the event of it having received a lot of punishments the network gets into an excited state with a high value of CCE and will frequently instantiate new cells.

[0068] Figure 7 shows a part of a network according to this embodiment, comprising two output nodes CR (correct response) and IR (incorrect response), an input layer including three input nodes a1, a2, a3 and two intermediate layer nodes CR1 and CR2, each connected to each of the three input nodes via respective weights (not shown), and to one of the output nodes. Several intermediate layer nodes may, in general, be connected to one output layer node. A group of intermediate nodes connected to a common output node is termed a 'cluster'. When an input signal held in short term memory contains the 3 features to which a1, a2 and a3 are responsive, the correct response should be for cell

CR1 to fire, triggering output cell CR, but assume that this is not the case, and that IR1 fires instead, triggering incorrect response IR. The supervisory program punishes the net and connections are varied according to the algorithm described above with reference to Figure 5. The net is then "told" that the output CR was correct, and further tests take place. The highest output of those nodes in the cluster connected to CR at the time of output-determination (in this example, since there is only one, it is that of CR1) is compared with that of the winning excitation which is that of IR1. The following logical step is then made:

If Excitation (CR1) < T.(Excitation IR1)), where $T = T(\text{CCE, other parameters}^*)$,

then a new cell, CR2, is instantiated and connected to CR, and appropriate connections made as shown in Figure 7b, so that CR2 gives the correct response, by causing CR to fire.

[0069] A further parameter*, termed the cell creation threshold (CCT) governs the level of excitement that is required to increase the size of the particular cluster to which it applied, and is related to the number of cells in that output cluster. If the cluster is large, it will be resistant to the creation of new cells. This means that the net will not instantiate cells very easily in a large cluster, so that there is a natural tendency towards economical use. In one example, the network described above was simulated on an IBM AT using muLISP. Two kinds of input were tested on the network: English sentences and phrases, and binary inputs.

English Phrases and Sentences

[0070] The network could learn to respond correctly to a large range of inputs in which deviation from the training inputs was permissible, and yet small but significant changes to the input were correctly interpreted. A simple example of a well trained net is shown in Figure 8. Sentences were taught so that the net learned the following set of responses:

"What is your name?:" - "Tania"

"What is your sister's name?" - "Gina"

"What is your younger sister's name?" - "Dolores"

"What is your job?" - "Gipsy"

[0071] It should be noted that with this easy set of inputs no clusters were formed, and the intermediate layer is not more than a set of single cells. Inputs with minor variations, such as "Whats your name"? or "What job do you do?" are correctly interpreted, of course.

[0072] Some variants of the net may allow punishment without subsequent teaching. In such cases it is possible, for example, to teach the net to respond to the input "What is your name" with, say, "Lizzie". A further period of training can teach the response "Elizabeth". If the net is then repeatedly asked to give its name, and punished when the reply "Elizabeth" is given, though without a correct response being taught, it ultimately reverts to its former name "Lizzie".

[0073] In summary the network training is able to store a vocabulary, and organise connections between items in the vocabulary in such a way that whilst unimportant variations in the input are correctly ignored, significant variation are correctly interpreted. The network responded quickly to any input even after it had learned moderately large vocabularies of input words (about four hundred).

Binary Inputs

XOR problem

[0074] Figure 9 shows a net which has been taught to solve the XOR problem. The four inputs L1 L2 O1 and O2 are used to generate the logical inputs 11 10 01 and 00, as described in the (LOOLLOLL) example above, so that for example the input (L1 O2) represents 10. (This representation can be made more neural with an extra layer of two input cells if numerical binary inputs are desired).

[0075] It will be noted that in this case two clusters of two cells each formed in the intermediate layer. In this case the XOR problem is solved by instantiating these four intermediate cells which prevent the hunting back and forward which must inevitably occur when only a single layer net is used, and precluding the necessity for the thousands of iterations required by a multilayer perceptron (MLP) with back-propagation (Rumelhart, Hinton & Williams 1986) to solve the XOR problem. It can be seen that a reduction in processing time has been obtained in exchange for an increase in memory. In practice three internal nodes are sometimes adequate to solve this problem, and at the other extreme with a bad training method as many as five are instantiated - with some obviously redundant. An algorithm for reclaiming memory wasted in this way is therefore desirable (but not essential). It can superficially appear that to

use so many cells to solve the XOR problem is counter productive, in that one might as well merely store the inputs and outputs as memories. In fact it is the ability of the net to decide for itself when features are worth storing in this way which gives it its power. As is shown below, some problems which involve much larger input patterns may have the XOR problem implicitly contained in them. In such cases although there may be many input nodes, only four intermediate nodes are formed, so that the solution seems to consist of a few memorised cases generalised to allow recognition of all patterns, which is a very efficient use of memory.

[0076] The two nets of Figures 8 and 9 were actually parts of the same net, which formed two completely independent subnets since there was no overlap between the sets of inputs and outputs. If it were wished to train the same net to solve some other logic function, then during training, instead of inputting only (L1 O2) for example, the input (XOR L1 L2) would be used, and this input associated with the XOR output. Then for example the OR function can be trained by using inputs like (OR L1 O2) with the appropriate output. This method allows the net to learn to handle inputs in different ways according to the wishes of the user.

Parity Problem

[0077] The network described could solve the third order parity problem when trained in a similar manner in somewhere between three and seven complete cycles of presentation of the eight input states.

Blob Finding in Video Images

[0078] Using inputs in the form of a table in which each column of Os and Is is regarded as a two-dimensional input pattern, the network was successfully trained to decide whether a cluster or blob of ones was generally north, south, east or west.

Implicit XOR

[0079] The network could also generalise its ability to solve the XOR problem, to solve pattern recognition problems where a disjoint region in a two-dimensional space is hidden in a sixteen dimensional space.

[0080] Noisy patterns were correctly identified. In this case it was established that four internal nodes were instantiated to cope with the complex training set, but all subsequent inputs excited one or other of these internal nodes, to produce correct output without further proliferation of internal nodes. This seems to offer great promise, since it avoids the enormous amount of processing required by, for example the MLP, and produces the correct answers by the extra, but economic, use of memory.

[0081] Connections within layers could be implemented in the invention, as there is evidence that permitting such interconnection is useful.

[0082] Thus, two particular ways of generalising the invention are possible, one in which intra-layer connections are permitted, and one in which the competition amongst outputs is performed by lateral inhibition rather than a global maximum function, so that more than one output cell may fire.

[0083] From the foregoing, it will be apparent that a hardware embodiment of the invention, in which each node in the intermediate and output layers is a simple parallel processor capable of summing inputs from the nodes of the layer below weighted by corresponding weights, is straightforward to design. Suitable hardware is discussed in the literature and may, as discussed in our international application published as WO 89/02134, and which is herein incorporated by this reference, be analogue, digital or a hybrid of the two. Also necessary during training is a more general purpose computing device, for example a microprocessor or DSP (digital signal processing) device, programmed to train the network by instantiating new nodes and adjusting weight values.

[0084] Alternatively, the invention may be realised as a computing device (a microprocessor, or preferably a DSP device) programmed to perform the parallel hardware functions (as parallel program branches) sequentially.

[0085] Finally, it is observed that networks trained according to the invention may be marketed, for specific applications such as phrasebook translations, as a trained device without the training means. Such devices will in general reflect the fact that they have been trained using the invention, and will hence not resemble known networks in their architecture; for example, nodes instantiated towards the end of training will be more sparsely connected.

REFERENCES

[0086]

[1] Allen, J. (1987) *Natural Language Understanding*. (1987) California: Benjamin/Cummings.

- [2] Alonso, J. A. and Schneider, T. (1989) *Machine Translation Technology: On the way to Market Introduction*. International Journal of Computer Applications in Technology. 2 pp 186-190.
- [3] Hutchinson, R.A. and Welsh, W.J. (1989) *Comparison of Neural Networks and Conventional Techniques for Feature Location in Facial Images*. First IEE International Conference on Artificial Neural Nets, London 16-18 October 1989.
- [4] Jogota, A. Jacobowitz, O (1989) *Knowledge Representation in Multilayered Hopfield Nets*. International Joint Conference on Neural Nets, Washington, June 19-22.
- [5] Lynch, M.R. and Rayner, P.J. (1989) *The Properties and Implementation of the Non-Linear Vector Space Connectionist Model*. First IEE International Conference on Artificial Neural Nets, London 16-18 October 1989.
- [6] McClelland, D.E. and Kawamoto, A.H. (1986) Mechanisms of Sentence Processing: Assigning Roles to Constituents of Sentences. *Parallel Distributed Processing*. 2 McClelland, J.L. and Rumelhart, D.E. (Eds.) Cambridge, Massachusetts: MIT Press.
- [7] Morton K, Coulston, M and Gerrihy, G (1989) *Translation English to French Limited Domain Translation using Dynamic Topology Net*. Report for British Telecom CONNEX project.
- [8] Rumelhart, D.E. and McClelland, J.L. (1986) *Parallel Distributed Processing*, Cambridge, Massachusetts: MIT Press.
- [9] Rumelhart, D.E. and McClelland, J.L. (1986) On Learning the Past Tenses of English Verbs. *Parallel Distributed Processing*. 2 McClelland, J.L. and Rumelhart, D.E. (Eds.) Cambridge, Massachusetts: MIT Press.
- [10] Rumelhart, D.E., Smolensky, P., McClelland, J.L. and Hinton, G.E. (1986) Schemata and Sequential Thought Processes in PDP models. *Parallel Distributed Processing*. 2 McClelland, J.L. and Rumelhart, D.E. (Eds.) Cambridge, Massachusetts: MIT Press.
- [11] Rumelhart, D.E. and Zipser D. (1986) Feature Discovery by Competitive Learning. *Parallel Distributed Processing*. 1 McClelland, J.L. and Rumelhart, D.E. (Eds.) Cambridge, Massachusetts: MIT Press.
- [12] Rumelhart, D.E., Hinton, G.E. and Williams, R.J. (1986) Learning Internal Representations. *Parallel Distributed Processing*. 2 McClelland, J.L. and Rumelhart, D.E. (Eds.) Cambridge, Massachusetts: MIT Press.
- [13] Saerens, M. and Soquet, A., (1989) *A Neural Controller* First IEE International Conference on Artificial Neural Nets, London 16-18 October 1989.
- [14] Sharkey, N.E. (1989) A PDP Learning Approach to Natural Language. *Neural Computing Architectures*. I. Alexander (Ed.), London: Kogan-Page.
- [15] Slocum, J. (1989) A Survey of Machine Translation: its History, Current Status, and Future Prospects" *Machine Translation Systems*. J. Slocum, (Ed.) Cambridge, England: Cambridge University Press.
- [16] Weber, S.H. (1989) *A Connectionist Model of Conceptual Representation*. International Joint Conference on Neural Nets, Washington, June 18-22.
- [17] Woodland, P.C. and Smyth, S.G. (1990) *An Experimental Comparison of Connectionist and Conventional Classification Systems on Natural Data*. To appear in Speech and Communication - Special Issue on Neurospeech.
- [18] Wyard, P., Nightingale, C. and Marsh R (1989) *A Higher Order Dynamic Topology Neural Net and its Application to natural language problems*. To be published.

Claims

1. A trainable artificial neural network comprising:

input layer means for receiving sets of data items, a data item being a data signal, which layer means is arranged to generate, in operation, node output signals each indicating the presence, if any, of a respective predetermined data item in a set of data items received by said input layer means; and
 5 output layer means arranged to receive, in operation, node output signals generated by the input layer means or signals derived from these node output signals and generate at least one net output signal, the or each net output signal depending upon a plurality of said node output signals weighted by respective weight values;

characterised in that the network further comprises:

10 storage means arranged to store, in operation, given data items, if any, which are present in a set of data items received by said input layer means but which do not correspond to said respective predetermined data items so to instantiate new nodes in said input layer means; and

15 input layer modifying means arranged to modify, in operation, the weight of the output corrections of the nodes of the input layer means so as to cause the input layer means thereafter to generate, in operation, node output signals each indicating the presence, if any, of a respective said given data item in a set of data items received by said input layer means.

2. A network according to Claim 1, in which each said set of data items is in the form of a sequence of corresponding data items and said given data items are subsequences of said data items in a said sequence.

20 3. A network according to Claim 2 in which said subsequences comprise data items occurring contiguously in a said sequence.

25 4. A network according to Claim 2 or Claim 3 in which said subsequences comprise fewer than a predetermined number of said data items.

30 5. A network according to Claim 2 or Claim 3 in which said subsequences are a subset of subsequences of said data items occurring in a said sequence, the probability of a given subsequence being included in the subset decreasing with increasing subsequence length.

35 6. A network according to Claim 1 further comprising input means for receiving sequences of data items, which input means is arranged to generate, in operation, said sets of data items in such manner that each generated set of data items corresponds to a received said sequence and comprises representations of respective data items in the corresponding sequence, each representation being indicative of both the identity of the data item it represents and the position of that data item in the corresponding sequence.

40 7. A network according to any preceding claim further comprising control means which is arranged to inhibit, in operation, the operation of said input layer modifying means in dependence upon the number of different data items which would currently result in the generation of a said node output signal by the input later means should these different data items be received by the input layer means, the dependence being such that relatively large degrees of inhibition will occur with relatively large values of said number.

45 8. A network according to any preceding claim, including further input layer modifying means arranged to modify, in operation, the input layer means in the event that a said predetermined data item has occurred relatively infrequently in sets of data items received by the input layer means, the modification being such as to cause the input layer means to fail to generate a node output signal indicating the presence of this predetermined data item in a set of data items subsequently received by the input layer means even if this predetermined data item should actually be present in the set.

50 9. A network as claimed in any preceding claim, implemented by means of a digital computer operating under stored program control.

10. A method of training an artificial neural network which comprises:

55 input layer means for receiving sets of data items, a data item being a data signal, which layer means is arranged to generate, in operation, node output signals each indicating the presence, if any, of a respective predetermined data item in a set of data items received by said input layer means; and
 output layer means arranged to receive, in operation, node output signals generated by the input layer means

or signals derived from these node output signals and generate at least one net output signal, the or each net output signal depending upon a plurality of said node output signals weighted by respective weight values;

characterised in that the method comprises the steps of:

inputting training sets of data items to the input layer means;
storing given data items which are present in a said training set but which do not correspond to said respective predetermined data items, so to instantiate new nodes in said input layer means,
detecting said net output signals; and, in the event that a predetermined criterion of success is not met,
modifying the weight of the output connections of the nodes of the input layer means so as to cause the input layer means thereafter to generate, in operation, a node output signal indicating the presence of at least one said given data item in a set of data items received by the input layer means if that given data item should be present in that set.

Patentansprüche

1. Trainierbares künstliches neuronales Netzwerk, das umfaßt:

- eine Eingangsschicht-Einrichtung zum Empfang von Komplexen von Datenangaben, wobei eine Datenangabe ein Datensignal ist, und die Eingangsschicht-Einrichtung so ausgelegt ist, daß sie im Betrieb Ausgangssignale der Knoten erzeugt, die das Vorliegen einer entsprechenden vorgegebenen Datenangabe, falls vorhanden, in einem Komplex von Datenangaben anzeigen, die von der Eingangsschicht-Einrichtung erhalten wurden, und
- eine Ausgangsschicht-Einrichtung, die so ausgelegt ist, daß sie im Betrieb Ausgangssignale der Knoten, die von der Eingangsschicht-Einrichtung erzeugt wurden, oder Signale erhält, die von diesen Ausgangssignalen der Knoten stammen, und zumindest ein Netzausgangssignal erzeugt, wobei das oder jedes Netzausgangssignal von mehreren Knotenausgangssignalen abhängig ist, die durch entsprechende Gewichtswerte gewichtet wurden,

dadurch gekennzeichnet, daß
das Netzwerk weiter umfaßt:

- eine Speichereinrichtung, die ausgelegt ist, im Betrieb die vorgegebenen Datenangaben, falls vorhanden, zu speichern, die in einem Komplex von Datenangaben vorliegen, die von der Eingangsschicht-Einrichtung erhalten wurden, die aber nicht den jeweiligen vorgegebenen Datenangaben entsprechen, so daß spontan neue Knoten in der Eingangsschicht-Einrichtung gebildet werden, und
- eine Eingangsschicht-Modifizierungseinrichtung, die ausgelegt ist, im Betrieb das Gewicht der Ausgangskorrekturen der Knoten der Eingangsschicht-Einrichtung so zu modifizieren, daß die Eingangsschicht-Einrichtung danach veranlaßt wird, im Betrieb Knotenausgangssignale zu erzeugen, die jeweils das Vorhandensein einer entsprechenden Datenangabe, falls vorhanden, in einem Komplex von Datenangaben anzeigen, die von der Eingangsschicht-Einrichtung erhalten wurden.

2. Netzwerk nach Anspruch 1, in dem jeder Komplex von Datenangaben die Form einer Sequenz entsprechender Datenangaben aufweist, und die vorgegebenen Datenangaben Subsequenzen der genannten Datenangaben in einer Sequenz sind.

3. Netzwerk nach Anspruch 2, in dem die Subsequenzen Datenangaben enthalten, die in einer Sequenz sequentiell auftreten.

4. Netzwerk nach Anspruch 2 oder 3, in dem die Subsequenzen weniger als eine vorgegebene Anzahl von Datenangaben aufweisen.

5. Netzwerk nach Anspruch 2 oder 3, in dem die Subsequenzen ein Unterkomplex der in einer Sequenz auftretenden Datenangaben sind, wobei die Wahrscheinlichkeit, daß eine gegebene Subsequenz im Unterkomplex enthalten ist, mit der zunehmenden Länge der Subsequenz abnimmt.

6. Netzwerk nach Anspruch 1, das ferner eine Eingabeeinrichtung zum Empfang von Sequenzen von Datenangaben

umfaßt, wobei die Eingabeeinrichtung so ausgelegt ist, daß sie im Betrieb die Komplexe der Datenangaben in der Weise erzeugt, daß jeder erzeugte Komplex von Datenangaben einer erhaltenen Sequenz entspricht und Darstellungen der jeweiligen Datenangaben in der entsprechenden Sequenz enthält, wobei jede Darstellung sowohl die Identität der Datenangabe, die sie darstellt, als auch die Position der Datenangabe in der entsprechenden Sequenz anzeigt.

7. Netzwerk nach einem der vorhergehenden Ansprüche, das ferner eine Steuereinrichtung umfaßt, die so ausgelegt ist, daß sie im Betrieb den Betrieb der Eingangsschicht-Modifizierungseinrichtung in Abhängigkeit von der Anzahl der verschiedenen Datenangaben sperrt, die gegenwärtig zur Erzeugung eines Knotenausgangssignals der Eingangsschicht-Einrichtung führen könnten, falls diese verschiedenen Datenangaben von der Eingangsschicht-Einrichtung empfangen werden würden, wobei die Abhängigkeit derart ist, daß bei relativ hohen Werten dieser Zahl ein relativ hoher Sperrungsgrad auftritt.

8. Netzwerk nach einem der vorhergehenden Ansprüche, das ferner eine Eingangsschicht-Modifizierungseinrichtung umfaßt, die so ausgelegt ist, daß sie im Betrieb die Eingangsschicht-Einrichtung in dem Fall modifiziert, daß eine vorgegebene Datenangabe in den Komplexen der Datenangaben, die von der Eingangsschicht-Einrichtung erhalten werden, relativ selten aufgetreten ist, wobei die Modifizierung derart ist, daß sie die Eingangsschicht-Einrichtung veranlaßt, es zu versäumen, ein Knotenausgangssignal zu erzeugen, das das Vorhandensein dieser vorgegebenen Datenangabe in einem Komplex von Datenangaben anzeigt, die von der Eingangsschicht-Einrichtung subsequent erhalten wurden, auch wenn diese vorgegebene Datenangabe in Wirklichkeit im Komplex vorhanden sein sollte.

9. Netzwerk nach einem der vorhergehenden Ansprüche, das mit Hilfe eines digitalen Computers realisiert ist, der von einem gespeicherten Programm gesteuert betrieben wird.

10. Verfahren zum Training eines künstlichen neuronalen Netzwerks, das umfaßt:

- eine Eingangsschicht-Einrichtung zum Empfang von Komplexen von Datenangaben, wobei eine Datenangabe ein Datensignal ist, und die Eingangsschicht-Einrichtung so ausgelegt ist, daß sie im Betrieb Knotenausgangssignale erzeugt, die das Vorliegen einer jeweiligen vorgegebenen Datenangabe, falls vorhanden, in einem Komplex von Datenangaben anzeigt, die von der Eingangsschicht-Einrichtung erhalten wurden, und
- eine Ausgangsschicht-Einrichtung, die so ausgelegt ist, daß sie im Betrieb Knotenausgangssignale, die von der Eingangsschicht-Einrichtung erzeugt wurden, oder Signale erhält, die von diesen Knotenausgangssignalen stammen, und zumindest ein Netzausgangssignal erzeugt, wobei das oder jedes Netzausgangssignal von mehreren Knotenausgangssignalen abhängig ist, die durch entsprechende Gewichtswerte gewichtet wurden,

dadurch gekennzeichnet, daß das Verfahren die folgenden Schritte umfaßt:

- Eingabe von Trainingskomplexen von Datenangaben in die Eingangsschicht-Einrichtung,
- Speichern der gegebenen Datenangaben, die in einem Trainingskomplex vorliegen, aber die nicht den jeweiligen vorgegebenen Datenangaben entsprechen, um so spontan neue Knoten in der Eingangsschicht-Einrichtung zu bilden, und
- Erfassung des Netzausgangssignals (der Netzausgangssignale) und, falls kein vorgegebenes Erfolgskriterium angetroffen wird,
- Modifizierung des Gewichts der Ausgangskorrekturen der Knoten der Eingangsschicht-Einrichtung in der Weise, daß die Eingangsschicht-Einrichtung danach veranlaßt wird, im Betrieb ein Knotenausgangssignal zu erzeugen, das das Vorliegen von mindestens einer gegebenen Datenangabe in einem Komplex von Datenangaben anzeigt, die von der Eingangsschicht-Einrichtung erhalten wurden, falls die gegebene Datenangabe in diesem Komplex vorliegt.

Revendications

1. Réseau neuronal artificiel pouvant subir un apprentissage comprenant :

un moyen de couche d'entrée destiné à recevoir des ensembles d'éléments de données, un élément de données étant un signal de données, lequel moyen de couche est agencé pour générer, en fonctionnement, des signaux de sortie de noeuds indiquant chacun la présence, si elle existe, d'un élément de données prédéterminé.

miné respectif dans un ensemble d'éléments de données reçu par ledit moyen de couche d'entrée, et un moyen de couche de sortie agencé pour recevoir, en fonctionnement, des signaux de sortie de noeuds générés par le moyen de couche d'entrée ou bien des signaux obtenus à partir de ces signaux de sortie de noeuds et générer au moins un signal de sortie de réseau, le signal de sortie de réseau ou chaque signal de sortie de réseau dépendant d'un certain nombre desdits signaux de sortie de noeuds pondérés par des valeurs de poids respectives,

caractérisé en ce que le réseau comprend en outre :

un moyen de mémorisation agencé pour mémoriser, en fonctionnement, des éléments de données fournis, s'ils existent, qui sont présents dans un ensemble d'éléments de données reçu par ledit moyen de couche d'entrée mais qui ne correspondent pas auxdits éléments de données prédéterminés respectifs, de façon à créer de nouveaux noeuds dans ledit moyen de couche d'entrée, et

un moyen de modification de couche d'entrée agencé pour modifier, en fonctionnement, le poids des connexions de sortie des noeuds du moyen de couche d'entrée, de façon à amener le moyen de couche d'entrée à générer ensuite, en fonctionnement, des signaux de sortie de noeuds indiquant chacun la présence, s'il en existe, d'un dit élément de données fourni respectif dans un ensemble d'éléments de données reçu par ledit moyen de couche d'entrée.

2. Réseau selon la revendication 1, dans lequel chaque dit ensemble d'éléments de données est sous forme d'une séquence d'éléments de données correspondants et lesdits éléments de données fournis sont des sous-séquences desdits éléments de données dans une dite séquence.
3. Réseau selon la revendication 2, dans lequel lesdites sous-séquences comprennent des éléments de données apparaissant de façon contiguë dans une dite séquence.
4. Réseau selon la revendication 2 ou la revendication 3, dans lequel lesdites sous-séquences comprennent moins qu'un nombre prédéterminé desdits éléments de données.
5. Réseau selon la revendication 2 ou la revendication 3, dans lequel lesdites sous-séquences constituent un sous-ensemble de sous-séquences desdits éléments de données apparaissant dans une dite séquence, la probabilité pour qu'une sous-séquence donnée soit incluse dans le sous-ensemble diminuant avec une longueur de sous-séquence croissante.
6. Réseau selon la revendication 1, comprenant en outre un moyen d'entrée destiné à recevoir des séquences d'éléments de données, lequel moyen d'entrée est agencé pour générer, en fonctionnement, lesdits ensembles d'éléments de données de telle manière que chaque ensemble généré des éléments de données correspond à une dite séquence reçue et comprend des représentations des éléments de données respectifs dans la séquence correspondante, chaque représentation étant indicative à la fois de l'identité de l'élément de données qu'il représente et de la position de cet élément de données dans la séquence correspondante.
7. Réseau selon l'une quelconque des revendications précédentes, comprenant en outre un moyen de commande qui est agencé pour inhiber, en fonctionnement, le fonctionnement dudit moyen de modification de la couche d'entrée suivant le nombre des éléments de données différents qui résulterait actuellement de la génération d'un dit signal de sortie de noeuds par le dernier moyen d'entrée, si ces différents éléments de données devaient être reçus par le moyen de couche d'entrée, la dépendance étant telle que des degrés relativement importants d'inhibition se produiraient avec des valeurs relativement grandes dudit nombre.
8. Réseau selon l'une quelconque des revendications précédentes, comprenant en outre un moyen de modification de couche d'entrée agencé pour modifier, en fonctionnement, le moyen de couche d'entrée dans le cas où un dit élément de données prédéterminé est apparu relativement peu fréquemment dans des ensembles d'éléments de données reçus par le moyen de couche d'entrée, la modification étant telle qu'elle amène le moyen de couche d'entrée à ne pas réussir à générer un signal de sortie de noeuds indiquant la présence de cet élément de données prédéterminé dans un ensemble d'éléments de données reçu ensuite par le moyen de couche d'entrée, même si cet élément de données prédéterminé devrait réellement être présent dans l'ensemble.
9. Réseau selon l'une quelconque des revendications précédentes, réalisé au moyen d'un ordinateur numérique fonctionnant sous la commande d'un programme mémorisé.

10. Procédé d'apprentissage d'un réseau neuronal artificiel qui comprend :

un moyen de couche d'entrée destiné à recevoir des ensembles d'éléments de données, un élément de données étant un signal de données, lequel moyen de couche est agencé pour générer, en fonctionnement, des signaux de sortie de noeuds indiquant chacun la présence, s'il en existe, d'un élément de données prédéterminé respectif dans un ensemble d'éléments de données reçu par ledit moyen de couche d'entrée, et un moyen de couche de sortie agencé pour recevoir, en fonctionnement, des signaux de sortie de noeuds générés par le moyen de couche d'entrée ou bien des signaux obtenus à partir de ces signaux de sortie de noeuds et générer au moins un signal de sortie de réseau, le signal de sortie de réseau ou chaque signal de sortie de réseau dépendant d'un certain nombre desdits signaux de sortie de noeuds pondérés par des valeurs de pondération respectives,

caractérisé en ce que le procédé comprend les étapes consistant à :

recevoir en entrée des ensembles d'apprentissage des éléments de données vers le moyen de couche d'entrée,

mémoriser des éléments de données fournis qui sont présents dans un dit ensemble d'apprentissage mais qui ne correspondent pas auxdits éléments de données prédéterminés respectifs, de façon à créer de nouveaux noeuds dans ledit moyen de couche d'entrée,

détecter ledit signal ou lesdits signaux de sortie de réseau, et, dans le cas où un critère de succès prédéterminé n'est pas satisfait,

modifier le poids des connexions de sortie des noeuds du moyen de couche d'entrée de façon à amener le moyen de couche d'entrée à générer ensuite, en fonctionnement, un signal de sortie de noeuds indiquant la présence d'au moins un dit élément de données fourni dans un ensemble d'éléments de données reçu par le dernier moyen d'entrée si cet élément de données fourni devrait être présent dans cet ensemble.

Fig.1




Fig.6

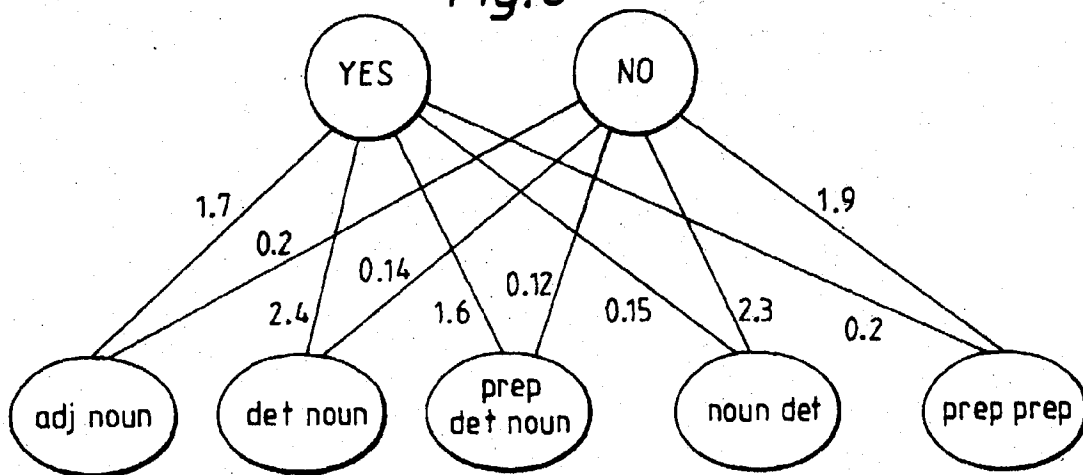


Fig 2a





Fig 2b





Fig 2c





Fig 2d





Fig. 3




Fig. 5


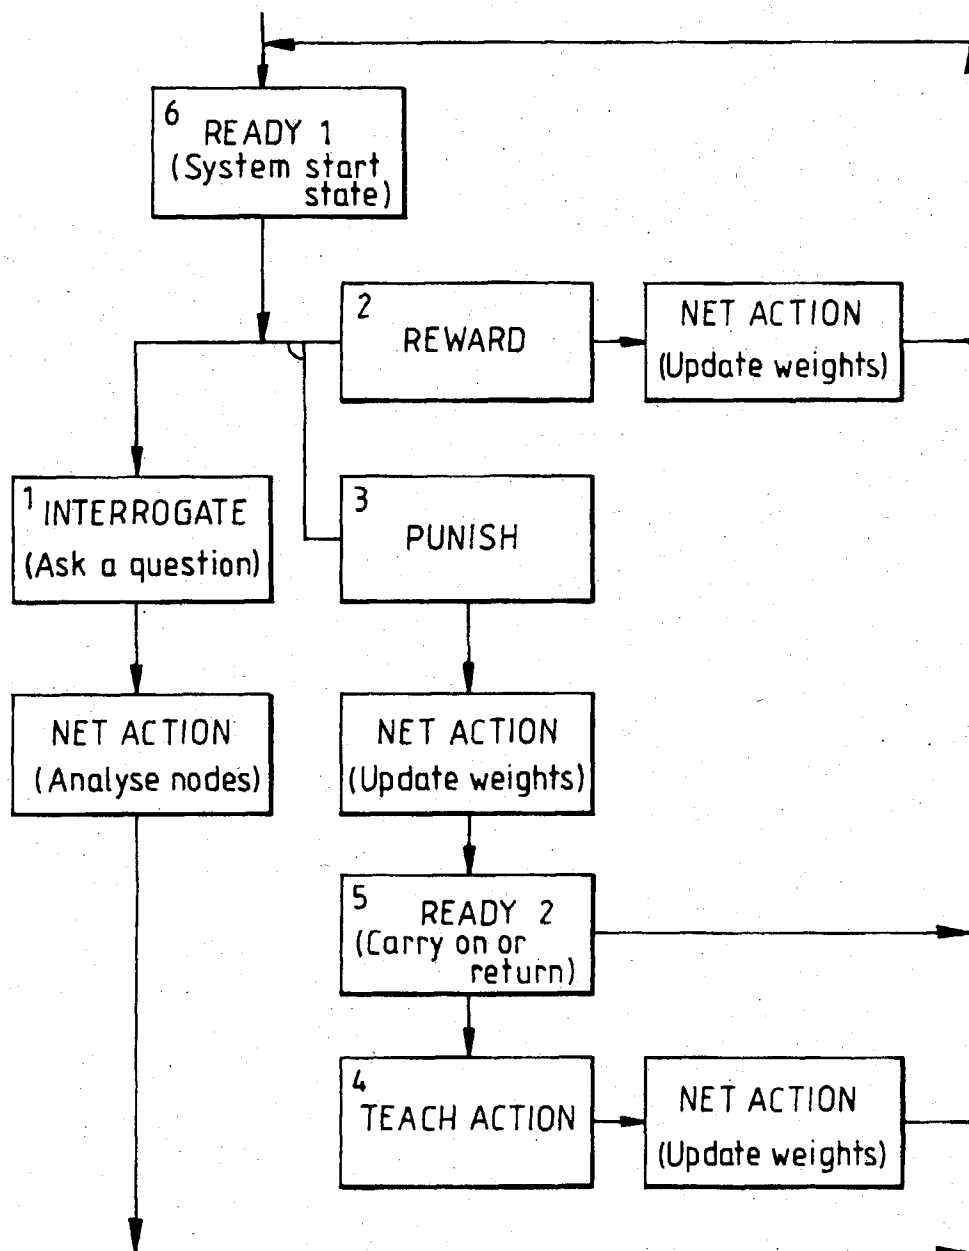


Fig. 4



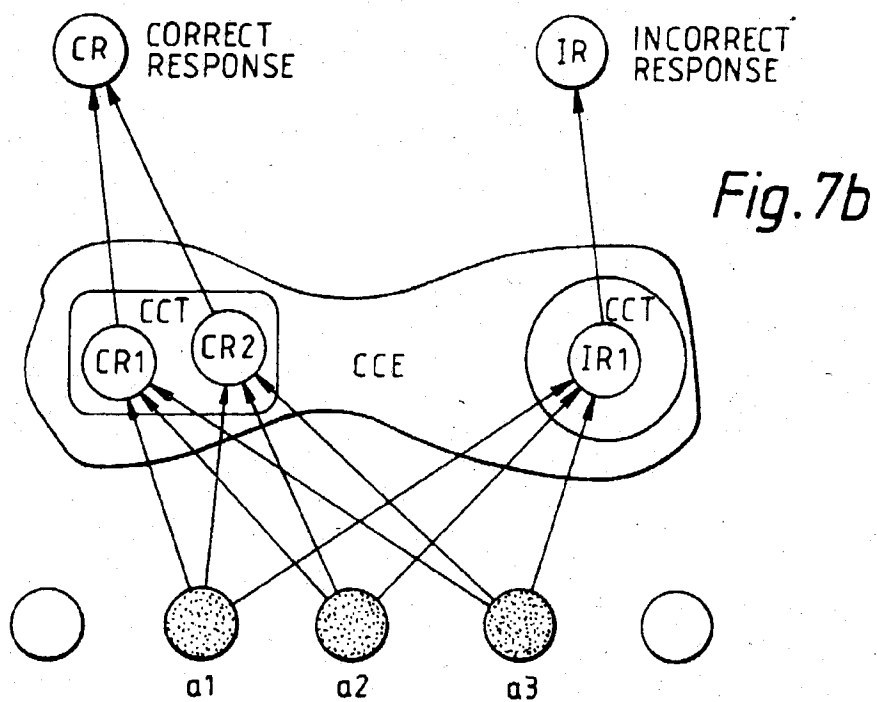
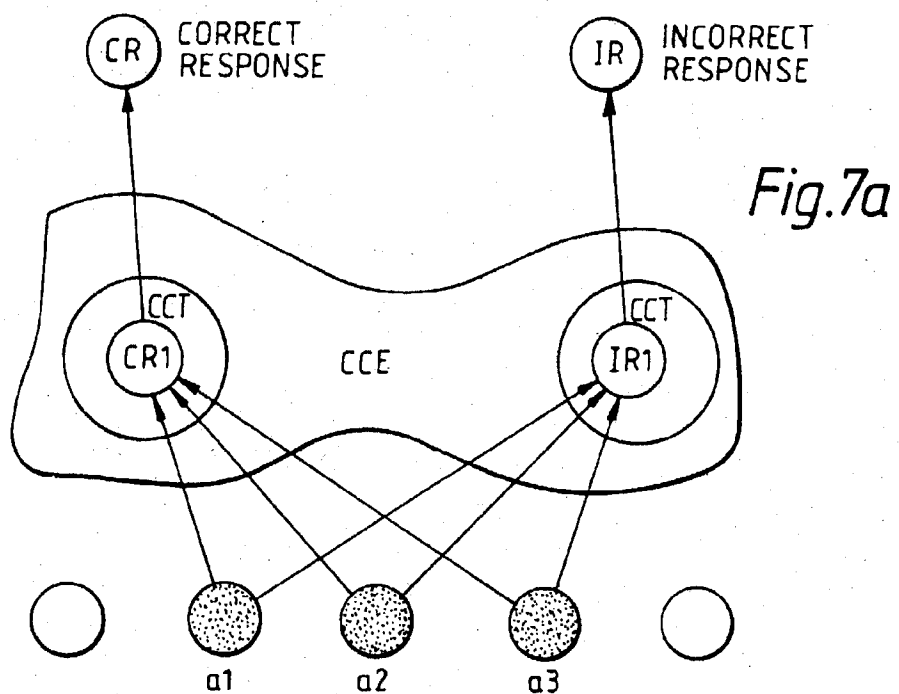


Fig 8

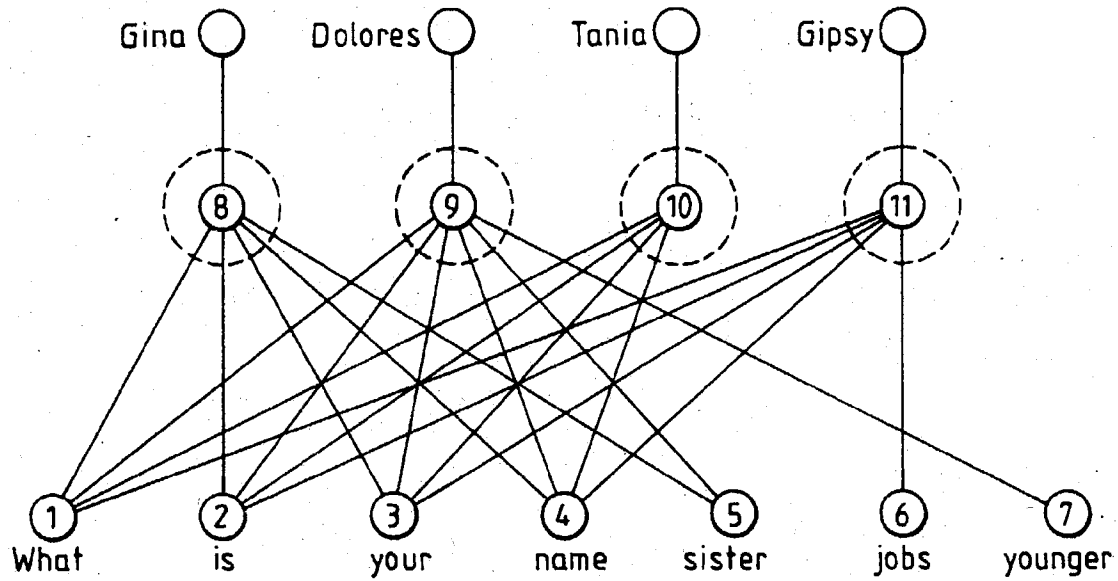


Fig. 9

